

Near Optimal Hamiltonian-Control and Learning via Chattering

PEEYUSH KUMAR, WOLF KOHN and ZELDA B. ZABINSKY, Atigeo, 800 Bellevue Way NE Suite 600, Bellevue, WA

Many applications require solving non-linear control problems that are classically not well behaved. This paper develops a simple and efficient chattering algorithm that learns near optimal decision policies through an open-loop feedback strategy. The optimal control problem reduces to a series of linear optimization programs that can be easily solved to recover a relaxed optimal trajectory. This algorithm is implemented on a real-time enterprise scheduling and control process.

Additional Key Words and Phrases: Relaxed optimal control problem, Variational optimization, Chattering approximation, Hamiltonian control, Continuous dynamical systems, Non-convex optimization

1. INTRODUCTION

Many decision problems involve solving a non-linear control problem that lacks traditional properties such as smooth controls and convex state spaces. This paper develops a chattering algorithm to solve non-linear control problems that are classically not well-behaved. The algorithm works with the Hamiltonian formulation for optimal control problems. This formulation provides many advantages as it can incorporate logic rules (the language of communication in enterprises), can facilitate planning and model adaptation (Hamiltonian describes the time evolution of a system), and scales to big data systems (Hamiltonians are additive for multi-body systems). In addition, this method is implemented using an open-loop feedback strategy that balances the trade-off between control adaptability and computational efficiency. We demonstrate this method on a real-time enterprise scheduling and control process.

Non-linear control problems are used in diverse applications but they remain challenging to solve computationally. One generally successful method is the State-Dependent Riccati Equation (SDRE) technique [Çimen, 2012]. While this method has been used in many applications, its major limitation is that it requires solving the Riccati equations multiple times. This is generally computationally inefficient. Another recent work in the field solves the non-linear optimal control problem using homotopy perturbation methods [Jajarmi et al., 2012]. This method is computationally efficient as compared to the SDRE technique. However, the approximation methods that are used lack formal analysis on error bounds. It is not well understood how these methods can generalize to ill-behaved, non-differentiable, non-convex systems.

Another well studied approach to solving non-linear control problems is dynamic programming and various approximate solutions enable better computational performance [Bertsekas, 2007; Powell, 2007]. The reinforcement learning and Bayesian learning literature has many algorithms using dynamic programming for Markov

This work is supported by Atigeo, 800 Bellevue Way NE Suite 600, Bellevue, WA 98004.

Author's addresses: Peeyush Kumar and Zelda B. Zabinsky, Industrial and Systems Engineering, University of Washington Seattle. Wolf Kohn, Atigeo, 800 Bellevue Way NE Suite 600, Bellevue, WA 98004.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© YYYY ACM. 0004-5411/YYYY/01-ARTA \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

systems [Sutton and Barto, 1998; Wiering and van Otterlo, 2014]. These methods are generally adaptive and can solve a wide variety of problems, although their success is largely limited to discrete Markovian systems or discretized approximations to continuous Markovian systems. Dynamic programming approaches do not scale well to high dimensional systems due to the curse-of-dimensionality for discrete problems. One common approach to handle large systems is to solve a tailored approximation to the system of interest. This paper introduces a generalizable, scalable approach to ill-behaved non-linear control problems.

Any control problem can be written as an optimization program with an objective to be optimized and with constraints on state variables. The goal is to compute a control $u(t)$ that minimizes the objective function:

$$J(t, x(t), u(t)) = \int_0^T g(t, x(t), u(t)) dt + \Psi(x(T)) \quad (1)$$

where $x(t)$ is the system state, which evolves according to the state equations

$$x(t) = x(0) + \int_0^t f(\tau, x(\tau), u(\tau)) d\tau, \quad x(0) = x_0, \quad t \in [0, T] \quad (2)$$

and the control $u(t) \in U$ is bounded. The objective function $J(t, x(t), u(t))$ is the total accumulated cost over the time horizon of the problem. It is composed of an operational running cost $g(t, x(t), u(t))$ and the operational terminal cost $\Psi(x(T))$. The function $f(\tau, x(\tau), u(\tau))$ is the continuous dynamics of the system which can be non-differentiable with respect to the state ($x(t)$) and the control variable ($u(t)$). The state set $X = \{x\}$ can in general be a non-convex set.

The Hamiltonian formulation for control problems was introduced by Pontryagin as part of his minimum principle [Athans and Falb, 2006]. In this formulation the optimal control problem is given in terms of the Hamiltonian $\mathcal{H}(t, x(t), p(t), u(t)) = g(t, x(t), u(t)) + p^T f(t, x(t), u(t))$ where $p(t)$ is the momentum (co-state). The necessary conditions for optimality are

$$\left(\frac{dx(t)}{dt} \right)^T = \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial p} \quad (3)$$

$$\left(\frac{dp(t)}{dt} \right)^T = - \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial x} \quad (4)$$

$$\mathcal{H}(t, x^*(t), p(t), u^*(t)) \leq \mathcal{H}(t, x(t), p(t), u(t)) \quad (5)$$

and optimal $u^*(t)$, and where $p(T) = \frac{\partial \Psi}{\partial x} \Big|_{x(T)}$ and $x(0) = x_0$.

Thus we seek an optimal trajectory $u^*(t)$ such that (3), (4) and (5) hold.

Solving (3), (4) and (5) directly is a two point boundary value problem because $x(0)$ and $p(T)$ are known. Unfortunately, there are no analytic solutions to solve such problems. Additionally, Pontryagin's necessary conditions of optimality are only valid for continuous systems.

The chattering approach approximates a solution to the optimal control problem defined by (3), (4) and (5) that are non-linear, non-convex and non-smooth. The chattering problem can also be applied to discrete systems by *continualizing* the discrete system. A good description of the advantages of approximating a discrete dynamical system using a continuous system model is discussed in the article by Nykamp, 2015]. Kohn and Nerode, 1999] provide a formal method to find a continuous approximation to discrete systems.

Section 2 describes the chattering approximation approach. Section 3 provides a numerical method, we call the chattering algorithm. Section 4 demonstrates this approach on a real time enterprise scheduling and control system.

2. CHATTERING APPROACH TO CONTROL

The chattering approach to optimal control problems is motivated through measure based controls. Kohn et al., 2010] provide an approach for optimization of a specific class of variational problems via chattering. This paper extends the chattering lemma [Kohn et al., 2010] to a chattering approach for generic optimal control problems. The main idea in this paper is to work directly with control Hamiltonians by converting (3), (4) and (5) into a sequence of standard variational problems. Each variational problem is defined over a small time interval Δt_i at time t_i , $i = 0, \dots, I - 1$, by partitioning $[0, T]$ into I intervals $[t_i, t_{i+1}]$ defined for each i with $\Delta t_i = t_{i+1} - t_i$ and such that $t_0 = 0$ and $t_I = T$.

The necessary conditions in (5) are relaxed by defining the optimization problem over a small time interval Δt_i , $\Delta t_i = t_{i+1} - t_i$, where the objective function is an accumulation of the Hamiltonian and the constraints are obtained by integrating (4) and (5) over the same interval $[t_i, t_{i+1}]$ for every i . The approach consists of finding a solution to the following sequence of variational problems.

$$\begin{aligned} & \inf_{u(t) \in U, t \in [t_0, t_1]} \int_{t_0}^{t_1} \mathcal{H}(t, x(t), p(t), u(t)) dt \\ & \text{subject to} \\ & x(t) = x_{t_0} + \int_{t_0}^t \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial p} dt \\ & p(t) = p_{t_1} - \int_{t_1}^t \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial x} dt \end{aligned} \tag{6}$$

$$\begin{aligned} & \inf_{u(t) \in U, t \in [t_1, t_2]} \int_{t_1}^{t_2} \mathcal{H}(t, x(t), p(t), u(t)) dt \\ & \text{subject to} \\ & x(t) = x_{t_1} + \int_{t_1}^t \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial p} dt \\ & p(t) = p_{t_2} - \int_{t_2}^t \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial x} dt \end{aligned} \tag{7}$$

⋮

$$\begin{aligned}
 & \inf_{u(t) \in U, t \in [t_{I-1}, t_I]} \int_{t_{I-1}}^{t_I} \mathcal{H}(t, x(t), p(t), u(t)) dt \\
 & \text{subject to} \\
 & x(t) = x_{t_{I-1}} + \int_{t_{I-1}}^t \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial p} dt \\
 & p(t) = p_{t_I} - \int_t^{t_I} \frac{\partial \mathcal{H}(t, x(t), p(t), u(t))}{\partial x} dt.
 \end{aligned} \tag{8}$$

with initial condition $x(0) = x_0$ and terminal transversality condition $p(T) = \frac{\partial \Psi}{\partial x} \Big|_{x(T)}$.

We apply the chattering lemma to convert the sequence of variational problems in (6), (7) and (8) over the interval $[0, T]$ to a sequence of linear optimization problems over time intervals $[t_i, t_{i+1}]$, $\forall i$.

LEMMA 2.1. (*Chattering Lemma c.f. [Kohn et al., 2010; Kohn et al., 2003; Roubicek, 1997]*) Let $c^1(t_i), \dots, c^K(t_i)$ be given integrable functions from interval $[t_i, t_{i+1}]$ into \mathbb{R}^n with compact support and let a constant $\epsilon > 0$ be given. Then there exists a chattering combination $g_{t_i}^\alpha$ defined for any $\alpha^{t_i} = (\alpha_1^{t_i}, \dots, \alpha_K^{t_i}) \in \mathbb{R}_+^K$, with $\sum_{k=1}^K \alpha_k^{t_i} = 1$ and $\alpha_k^{t_i} \geq 0$ such that

$$\max_{t \in [t_i, t_{i+1}]} \left| \int_{t_i}^t \left(g_{t_i}^\alpha(\tau) - \sum_{k=1}^K \alpha_k^{t_i} c_k^{t_i} \right) d\tau \right| < \epsilon$$

and

$$\text{meas}\{t : g_{t_i}^\alpha(t) \neq g^\beta(t)\} \rightarrow 0, \text{ as } \beta \rightarrow \alpha^{t_i}, \forall \beta$$




Fig. 1. Partitioning Schema: x_{t_i}, p_{t_i} are instantiated at the beginning of the interval $[t_i, t_{i+1}]$ while the chattering function $c_{t_i}^k$ and measure $\alpha_{t_i}^k$ are defined over the interval $[t_i, t_{i+1}]$.

Figure 1 provides a graphical representation of the partitioning schema, where $x(t_i) = x_{t_i}, p(t_i) = p_{t_i}$ are instantiated at the beginning of the interval $[t_i, t_{i+1}]$ while the chattering function $c_k^{t_i}$ and measure $\alpha_k^{t_i}$ are defined over the interval $[t_i, t_{i+1}]$. The chattering levels $c_k^{t_i}$ can be any function over the interval $[t_i, t_{i+1}]$ but we assume them to be constant for the purpose of this paper.

Figure 2 gives a graphical representation of the chattering in the interval $[t_i, t_{i+1}]$. The chattering measure $\alpha_k^{t_i} = \alpha^{t_i}(s)$, where $s = [k, k+1]$ for some k , can be interpreted as a probability that in an infinitesimally small neighborhood Δ_t around t , we can find $c_k^{t_i}(\tau) \in s$ for all $\tau \in \Delta_t$. For illustration $k = 1, \dots, 5, s = [3, 4], t = t_i$ and $\Delta_t = t_{i+1} - t_i$ in Figure 2.

The following theorem enables us to derive the chattering formulation for the optimization problem in (6), (7) and (8).




Fig. 2. Graphical description of chattering

THEOREM 2.2. Consider the sequence of variational problems in (6), (7) and (8), $i = \{1, \dots, I-1\}$. For each $[t_i, t_{i+1}]$, define levels $c_k^{t_i}, k = 1, \dots, K$ that discretize the range of $u \in U$ such that they form an ordered bounded set. Then there exist $\alpha_k^{t_i}, k = 1, \dots, K; i = 0, \dots, I-1$ and $\sum_{k=1}^K \alpha_k^{t_i} = 1, \forall i$ such that the chattering combination $g_{t_i}^\alpha$ for $t \in \Delta_{t_i}^{t_i}$ closely approximates, in a Lebesgue norm sense, a solution $u^*(t)$ to the variational problem, and, moreover,

$$\int_{t_i}^{t_i + \Delta_{t_i}} \mathcal{H}(t, x(t), p(t), u(t)) dt \approx \sum_{k=1}^K \mathcal{H}(t_i, x_{t_i}, p_{t_i}, c_k^{t_i}) \alpha_k^{t_i} \Delta_{t_i}.$$

where $t_{i+1} = t_i + \Delta_{t_i}$ and $x(t_i) = x_{t_i}$ and $p(t_i) = p_{t_i}$

Proof. Let us consider a sequence $c_k^{t_i}, k = 1, \dots, K$ in U . Define a Young's measure $\alpha(t, S)$, for $S \subseteq U$, as the limit of the sequence as $k \rightarrow \infty$. We associate with $\alpha(t_i, s), \forall s \in S$ a constant measure $\alpha^{t_i}(s)$, such that $\alpha(t_i, s) = \alpha^{t_i}(s)$ for $t \in [t_i, t_{i+1}]$ and $\alpha^{t_i}(s) = \sum_{k=1}^K \alpha_k^{t_i} \delta(c_k^{t_i})$, where the levels $c_k^{t_i}, k = 1, \dots, K$ discretize the range of $u \in U, \sum_{k=1}^K \alpha_k^{t_i} = 1$ and $\alpha_k^{t_i} \geq 0, \forall i$ and $\delta(c_k^{t_i})$ are Dirac measures concentrated on $c_k^{t_i}$. Then

$$u(t) = \int_U s \alpha^{t_i}(ds)$$

where the integral is written as (with a slight abuse of notation)

$$\int_U s \sum_{k=1}^K \alpha_k^{t_i} \delta(c_k^{t_i}) = \int_U \alpha^{t_i}(s) \sum_{k=1}^K s \delta(c_k^{t_i}) = \sum_{k=1}^K \alpha_k^{t_i} c_k^{t_i}.$$

This implies that the last sum gives a close approximation to the solution of the relaxed problem in (6), (7) and (8) which is a close approximation to the solution of the original variational problem (1) and (2). Applying the chattering lemma the last sum above is arbitrarily closely approximated by the chattering combination for all i .

Therefore, by the properties of a Lebesgue integral and \mathcal{H} continuous in (x, p) and piecewise continuous in u , for any $u \in U$,

$$\int_{t_i}^{t_i + \Delta_{t_i}} \mathcal{H}(t, x(t), p(t), u(t)) dt \approx \sum_k \mathcal{H}(t_i, x_{t_i}, p_{t_i}, c_k^{t_i}) \alpha_k^{t_i} \Delta_{t_i}$$

for small enough Δ_{t_i} . \square

The chattering approximation of $u(t)$ on t yields

$$\int_{t_i}^{t_{i+1}} u(s) ds \approx \sum_{k=1} \alpha_k^{t_i} c_k^{t_i} \quad (9)$$

for $i = 0, \dots, I - 1$, which also implies that the chattering approximation to equations (3) and (4) (in the integral form) can be written as

$$x_{t_{i+1}} = x_{t_i} + \sum_{k=1}^K \alpha_k^{t_i} \frac{\partial \mathcal{H}(t, x(t), p(t), c_k^t)}{\partial p} \Big|_{t=t_i} \Delta_{t_i} \quad (10)$$

$$p_{t_{i+1}} = p_{t_i} - \sum_{k=1}^K \alpha_k^{t_i} \frac{\partial \mathcal{H}(t, x(t), p(t), c_k^t)}{\partial x} \Big|_{t=t_i} \Delta_{t_i} \quad (11)$$

Note that this chattering approximation can in general be written for any variable whose evolution depends on the variable $u(t)$.

3. LINEAR OPTIMIZATION PROBLEM AND NUMERICAL METHOD

Now we use the chattering approximation in Section 2 to formulate a sequence of linear optimization programs where we optimize over α for each time interval indexed by i .

The linear optimization program based on the approximations in Theorem 2.2 is

$$\min_{\alpha_k^{t_i}} \sum_k \mathcal{H}(t_i, x_{t_i}, p_{t_i}, c_k^{t_i}) \alpha_k^{t_i} \quad (12)$$

subject to

$$\sum_k \alpha_k^{t_i} = 1 \quad (13)$$

$$0 \leq \alpha_k^{t_i} \leq 1, \forall k \quad (14)$$

given x_{t_i}, p_{t_i} and known $c_k^{t_i}$, for each $i = 0, \dots, I - 1$. This is a relaxed knapsack problem which has an analytic solution. Hence, there is no computational cost in solving this linear program.

The control problem can now be solved using the above chattering formulation and the known initial condition on the state ($x(0) = x_0$) and a good estimate of the initial value of the co-state. In practical problems it is almost impossible to know apriori a good estimate of the initial co-state. In order to find a good estimate

4 RESULTS

of the initial condition we use a similar procedure as in the numerical method of *variation of extremals* (see [Kirk, 2004]) but while solving the chattering optimization program described above. This approach converts the two point boundary value problem into an initial value problem, which is much easier to solve numerically.

The algorithm starts by guessing the initial value of the co-state ($p(0) = p_0$) and iteratively corrects on this guess based on the final observation $p(T)$. This adaptation is done by considering a perturbation on the initial condition of the co-state. Given (x_{t_i}, p_{t_i}) and the chattering levels $c_{t_i}^k$, the optimization program given by (12), (13) and (14) is solved to obtain the optimal chattering measure $\alpha_{t_i}^k$. The chattering levels and the chattering measure along with equations (10) and (11) are used to obtain the trajectory (x_{t_i}, p_{t_i}) .

In order to propagate (x_{t_i}, p_{t_i}) starting at (x_0, p_0) , the algorithm defines the chattering levels $c_k^{t_i}$ for each time interval from index i to $i+1$, such that they cover the control set U and that any bounded constraints on $x_{t_{i+1}}, \forall i$ are not violated. Additionally, in order to find a suitable bound for the levels, $c_{max}^{t_i} = \max_k c_k^{t_i}, \forall i$ and $c_{min} = \min_k c_k^{t_i}$, the following equation has to be satisfied

$$x_{min} \leq x_{t_i} + f(x_{t_i}, u_{t_i})\Delta_{t_i} \leq x_{max}. \quad (15)$$

Given x_{min} and x_{max} , bounds on chattering levels can be determined numerically from (15). This will ensure that the chattering levels cover the set U and the state constraints are satisfied.

The algorithm also propagates a perturbed trajectory that starts at the initial state $x_{t_i}^\delta = x_0$ and the perturbed initial co-state $p_0^\delta = p_0 + \delta p$, for small δp . The corresponding trajectory for state and co-state, starting at (x_0, p_0^δ) at index i is given by $(x_{t_i}^\delta, p_{t_i}^\delta)$.

The algorithm then corrects the initial estimate of co-state by adapting the guess estimate based on deviation of the propagated co-state ($p(T)$) starting from p_0 from the actual final value of the co-state $\left. \frac{\partial \Psi}{\partial x} \right|_{x(T)}$. This correction is done using

$$p_0 \leftarrow p_0 + \gamma \left(\left[\frac{\partial^2 \Psi}{\partial x^2}(x(T)) \right] P_x(T) - P_p(T) \right)^{-1} \left(p_T - \frac{\partial \Psi}{\partial x}(x(T)) \right)$$

where γ is a step parameter, $P_x(t_i) = \left. \frac{\partial x}{\partial p_0} \right|_{x=x_{t_i}, p=p_{t_i}}$ and $P_p(t_i) = \left. \frac{\partial p}{\partial p_0} \right|_{x=x_{t_i}, p=p_{t_i}}$. The partial derivative can be numerically estimated by using a finite difference method. This process is repeated until convergence, i.e., $\|p_T - \frac{\partial \Psi}{\partial x}(x(T))\| < \epsilon$. At convergence the solution obtained for x_{t_i} is the optimal control trajectory and the solution obtained for $u_{t_i} = \sum_{k=1} \alpha_k^{t_i} c_k^{t_i}$ is the optimal control signal. The triple $\{t_i, x_{t_i}, u_{t_i}\}$ defines a control law (policy).

The procedure described above gives a simple routine for solving the optimal control problem, which is given by Algorithm 1.

4. RESULTS

This section presents two applications that demonstrate the chattering control method described in this paper. The first example is used as a proof-of-concept. We use a standard linear quadratic oscillator to compare the solution from the chattering algorithm (Alg. 1) described in this paper with the known analytic solution. The second demon-

ALGORITHM 1: Algorithm for Hamiltonian-Control via Chattering

Input: Historical data containing the sequence $\{t, x(t), u(t), x'(t)\}$, where $x'(t)$ is the next state obtained on applying control $u(t)$ at time t while in state $x(t)$. The criterion $(g(t, x(t), u(t))$ and $\Psi(x)$) that needs to be optimized.

Result: A sequence of (near) optimal control law $\{t_i, x_{t_i}, u_{t_i}\}$ for $i = 0, \dots, I$.

begin

Estimate the Hamiltonian $\mathcal{H}(t, x(t), p(t), u(t))$ from the given data or estimate the dynamics f (using generalized Kalman filter or any other equivalent method) and compute

$\mathcal{H} = g(t, x(t), u(t)) + p^T f(t, x, u);$

If the system is discrete use a continuous approximation (via *continualization*);

Initialize $i = 0$, $t_i = 0$, $x(0) = x_0$ and guess $p(0) = p_0$ and $p_0^\delta = p_0 + \delta p$ for a fixed δp ;

flag = False;

while flag is False or iter \leq maximumIterations **do**

while $i < I$ **do**

Measure x_{t_i} ;

Determine chattering levels $c_k^{t_i}$ such that they satisfy bounded constraints on $u(t)$;

Solve the optimization program (12), (13) and (14) twice, once with (x_{t_i}, p_{t_i}) and

second one with $(x_{t_i}^\delta, p_{t_i}^\delta)$;

Compute $u_{t_i} = \sum_{k=1}^K \alpha_k^{t_i} c_k^{t_i}$;

Compute $p_{t_{i+1}}, p_{t_{i+1}}^\delta, x_{t_{i+1}}^\delta$ using (10) and (11);

Store $x_{t_i}, p_{t_i}, u_{t_i}$

end

if $\|p_T - \frac{\partial \Psi}{\partial x}(x(T))\| < \epsilon$ **then**

flag = True

else

Compute $P_x(T) = \frac{x_T^\delta - x_T}{\delta p}$, $P_p(T) = \frac{p_T^\delta - p_T}{\delta p}$;

$p_0 \leftarrow p_0 + \gamma \left(\left[\frac{\partial^2 \Psi}{\partial x^2}(x(T)) \right] P_x(T) - P_p(T) \right)^{-1} (p_T - \frac{\partial \Psi}{\partial x}(x(T)))$

end

end

end

stration is on a real-time enterprise scheduling and control process. This enterprise problem was first introduced by Kohn et al., 2005]. This problem is a hybrid of continuous and discrete systems that is highly non-linear, non-convex and non-smooth.

4.1. Linear Quadratic Controller

A linear quadratic controller is a system with quadratic operational cost that is constrained to move on a linear manifold.

The control problem is

$$\min_{u(t)} \int_0^1 x(t)^2 + u(t)^2 dt$$

s.t.

$$\dot{x}(t) = x(t) + u(t)$$

with initial condition $x(0) = 10$.

The Hamiltonian for the linear quadratic controller is given by

$$\mathcal{H}(x, p, u) = x^2 + u^2 + p(x + u)$$




Fig. 3. Linear Quadratic Controller

Figure 3 compares the state trajectory and the near-optimal control obtained using the chattering method vs the state trajectory and optimal control obtained through the analytic solution.

As observed the chattering method provides a good approximation to the optimal state trajectory, see Fig. 3(a). The *chattering* artifact can be observed in Fig. 3(b), where the controller chatters on discretized control levels.

4.2. Real Time Enterprise Scheduling and Control Problem

The second example uses a modified version of the problem defined by Kohn et al., 2005]. Consider a food distribution system with 14 suppliers, one warehouse, and three customers (graphically illustrated in Figure 4).

There are five types of food items to be distributed, each type is supplied by one or more suppliers and all items are perishable. The demands by the three customers are given in Figure 5. The customers are given an importance factor of: (i) 1.0 for most preferable, whose orders always need to be satisfied; ii) 0.40 for average customers whose orders need to be satisfied completely in 40% of the cases; and (iii) 0.25 for least preferable customers whose orders need to be satisfied completely in 25% of the cases. The problem data includes: (i) a list of customers with ranking and demands for each item type; (ii) a list of suppliers with the lead times, fixed and variable costs of delivery and min and max order quantities; and (iii) a list of item types with warehouse carry-over cost per item type, penalties per unit time for not satisfying demands, and days of life before perishing. Tables I and II provide the values of these external parameters






Fig. 4. A distribution system including suppliers, warehouses, and customers.


of the grocer.




(a) Apples




(b) Oranges



(c) Banana



(d) Tea



(e) Olives

Fig. 5. Demand

Table I. Grocer's External Parameters

| Item j | Name | InvCarrCost γ^j | Penalty δ^j | Customer i | Importance w_i |
|----------|--------|------------------------|--------------------|--------------|------------------|
| 0 | Apple | 100 | 10 | 1 | 1.0 |
| 1 | Orange | 150 | 25 | 2 | 0.4 |
| 2 | Banana | 200 | 39 | 3 | 0.25 |
| 3 | Tea | 50 | 30 | | |
| 4 | Olive | 65 | 25 | | |

Table II. Suppliers' External Parameters

| Item j | Supplier k | UnitCost α_k^j | FixCost β_k^j | MinQty $\mu_{k \min}^j$ | MaxQty $\mu_{k \max}^j$ |
|----------|---------------|-----------------------|---------------------|-------------------------|-------------------------|
| 0 | New Hampshire | 20 | 10 | 7 | 14 |
| 0 | Colorado | 25 | 7 | 4 | 11 |
| 1 | Florida | 50 | 10 | 5 | 20 |
| 1 | California | 70 | 5 | 4 | 13 |
| 2 | Costa Rica | 20 | 15 | 8 | 13 |
| 2 | Italy | 30 | 20 | 2 | 13 |
| 2 | India | 15 | 25 | 6 | 25 |
| 3 | India | 12 | 25 | 2 | 16 |
| 3 | Sri Lanka | 11 | 25 | 9 | 26 |
| 3 | England | 20 | 15 | 6 | 14 |
| 3 | Market | 23 | 20 | 2 | 18 |
| 4 | Greece | 20 | 15 | 15 | 17 |
| 4 | Italy | 25 | 12 | 10 | 22 |
| 4 | Market | 30 | 18 | 11 | 14 |

The problem is to allocate the requested quantities $v_i^j(t)$ of items (j) to customers (i) by ordering $\mu_k^j(t)$ of items j from appropriate supplier k while minimizing the total incremental cost. The dynamics model consists of two elements (i) market conservation; and (ii) inventory conservation.

— The market conservation dynamics are given by

$$\dot{Z}_i^j(t) = -Z_i^j(t) + \Theta_i^j(t) - v_i^j(t), \quad i = 1, \dots, 3, \quad j = 1, \dots, 5$$

where, $Z_i^j(t)$ is the unsatisfied demand of product j to customer i at time t , $\Theta_i^j(t)$ is the continulized demand of product j to customer i at time t and $v_i^j(t)$ is the continualized delivery action of product j to customer i at time t .

— The inventory conservation dynamics are given by

$$\dot{X}^j(t) = -X^j(t) + \sum_{k \in S_j} s_k^j(t) + \hat{\mu}^j(t) - p^j(t), \quad j = 1, \dots, 5$$

where, $X^j(t)$ is the inventory of product j at time t , $\sum_{k \in S_j} s_k^j(t)$ is the total rate of supply of product j from suppliers k at time t , $\hat{\mu}^j(t) = \sum_{k \in S_j} \mu_k^j(t)$ is the cumulative ordering action of product j from supplier $k \in S_j$ at time t and $p^j(t)$ is the cumulative rate of delivery of product j to customers i .

The cost function to be minimized is given by

$$J(t) = - \sum_{i=1}^3 \sum_{j=1}^5 v_i^j(t) r_i^j(t) + \sum_{j=1}^5 (\alpha^j \hat{\mu}^j(t) + \beta^j) + \sum_{j=1}^5 \gamma^j X^j(t) \sum_{i=1}^3 \sum_{j=1}^5 \bar{w}_i^j Z_i^j(t)$$

where, $r_i^j(t)$ is the unit revenue obtained by selling product j to customer i at time t , γ^j is the inventory cost of carrying product j per unit time, $\alpha^j = \sum_k \alpha_k^j$ is the

envelop unit cost of ordering product j , $\beta^j = \sum_k \beta_k^j$ is the envelop fixed cost of ordering product j , $\bar{w}_i^j = \frac{w_i \delta^j}{\sum_{i=1}^3 \sum_{j=1}^5 w_i \delta^j}$ is the cost for not satisfying demand, δ^j is the penalty for not satisfying demand for product j and w_i is the relative weight for not satisfying the demand for product j . We use the function $J(t) |J(t)|$ as the objective function to be minimized (the reason we chose this function is because it provides stable and bounded solutions). For this system the state variable is $x(t) = [X(t), Z(t)]^T$. The control variable is $u(t) = [\mu(t), v(t)]^T$. The state space is continuous with dimension $d_X = 20$ and the control space is also continuous with dimension $d_U = 29$.

Figure 6 illustrates the delivery amount of product j to customer i ($v_i^j(t)$) and the inventory levels $X^j(t)$, computed using the chattering method described in this paper. The algorithm converged in less than 100 iterations. The total computer time was approximately 20 minutes on a single core CPU with 2.5GHz and 8GB RAM.

Conventional optimal control methods that use dynamic programming would be intractable on this system. This is due to the reason that dynamic programming (DP) works on a discretized formulation of the control problem. Even though the dimensionality of the continuous system is not very high, the discretization results in exponential number of states. Feature extraction, based on tile coding of the state space with unit size, with state space (X) bounded by $0 \leq x \leq [100, \dots, 100]$, and tile coding of the control space with unit size, with control space (U) bounded by the values given in Table II for order quantities μ and $0 \leq v \leq [20, \dots, 20]$ for delivery quantities, would give a discretized system of net dimension $X^{d_X} \times X^{d_X} \times \{\times U\}_{d_U} \approx 100^{40} \times 20^{29}$. Since not all controls are simultaneously admissible, this reduces to $\sim 10^{102}$. Hence, DP algorithms run into the curse-of-dimensionality. Additionally, parametric and approximate methods would be difficult to condition due to the nature of the problem. Conventional methods based on the Pontryagin's principle are unable to solve this problem due to the problem being non-smooth and non-convex, while the Hamiltonian chattering method achieves a good solution efficiently.

5. DISCUSSION AND CONCLUSION

This paper develops a chattering approach to Hamiltonian control problems that are not necessarily convex, linear or smooth. This method is observed to be robust with respect to the structure of the underlying problems as it does not assume any conditions on the state-control space of the system. Additionally, it provides a good approximation for non-smooth controls within a dense subset of the space. Computationally, the optimal control problem reduces to a sequence of linear optimization programs of size equal to the number of chattering levels which is of the order of the dimension of the state variable. The linear optimization programs are relaxed knapsack problems which can be solved analytically. In practice, we find that the algorithm converges in a very few number of iterations to the optimal solution.

This method also provides a computationally efficient way of solving continuous space-time control systems. The chattering approach intrinsically discretizes the problem which enables a numerical solution. Note that the discretization, using this method, is of the order of the dimension of the state space, as opposed to being exponential in conventional canonical discretization approaches. The control problem intrinsically is a 2-point boundary value problem, which unfortunately, is very challenging to solve. This method uses a variation of extremals to convert the two-point boundary value control problem into an initial value problem, which can be solved numerically.

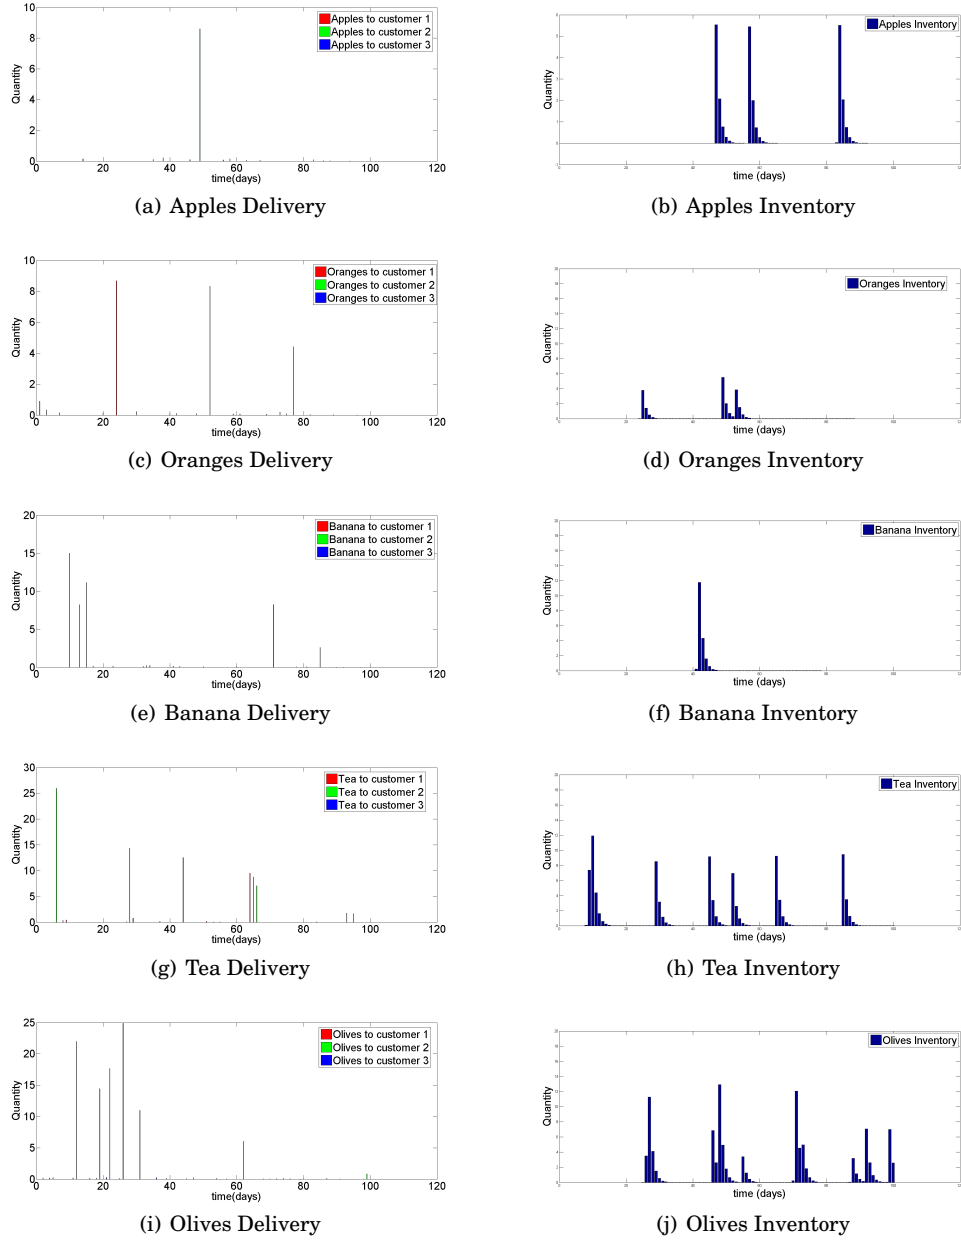


Fig. 6. Delivery and Inventory

For future research, we are focusing on finding more efficient and parallelizable solutions to control problems by exploring the theory of parallel transport in manifolds. The theory of parallel transport is well studied in the differential geometry community. We aim to use these results to efficiently find the neighborhood around the optimal trajectory.

ACKNOWLEDGMENTS

The authors would like to thank Atigeo Corporation for the generous support to carry out this research.

REFERENCES

- Athans, M. and Falb, P. (2006). *Optimal Control: An Introduction to the Theory and Its Applications*. Dover Books on Engineering Series. Dover Publications.
- Bertsekas, D. P. (2007). *Dynamic Programming and Optimal Control, Vol. II*. Athena Scientific, 3rd edition.
- Çimen, T. (2012). Survey of state-dependent riccati equation in nonlinear optimal feedback control synthesis. *Journal of Guidance, Control, and Dynamics*, 35(4):1025–1047.
- Jajarmi, A., Pariz, N., Kamyad, A. V., and Effati, S. (2012). A highly computational efficient method to solve nonlinear optimal control problems. *Scientia Iranica*, 19(3):759 – 766.
- Kirk, D. (2004). *Optimal Control Theory: An Introduction*. Dover Books on Electrical Engineering Series. Dover Publications.
- Kohn, W., Brayman, V., Cholewinski, P., and Nerode, A. (2003). Control in hybrid systems. *International Journal of Hybrid Systems*, 3(2):3.
- Kohn, W., Brayman, V., and Littleton, J. (2005). Repair-control of enterprise systems using rfid sensory data. *IEEE Transactions*, 37(4):281–290.
- Kohn, W., Brayman, V., Zabinsky, Z. B., and Shen, Y. (2010). Variational optimization via chattering. *Non-linear Analysis: Hybrid Systems*, 4(4):736 – 742.
- Kohn, W. and Nerode, A. (1999). Multiple-agent hybrid control architecture for intelligent real-time control of distributed nonlinear processes. US Patent 5,963,447.
- Nykamp, D. (2015). From discrete dynamical systems to continuous dynamical systems. From Math Insight.
- Powell, W. B. (2007). *Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics)*. Wiley-Interscience.
- Roubicek, T. (1997). *Relaxation in Optimization Theory and Variational Calculus*. De Gruyter Series in Nonlinear Analysis and Applications. De Gruyter.
- Sutton, R. S. and Barto, A. G. (1998). *Introduction to Reinforcement Learning*. MIT Press, Cambridge, MA, USA, 1st edition.
- Wiering, M. and van Otterlo, M. (2014). *Reinforcement Learning: State-of-the-Art*. Springer Publishing Company, Incorporated.